4

AD-A219 991

GL-TR-89-0174
INSTRUMENTATION PAPERS, NO. 338

THE ZENITH Z-248 AS A SCIENTIFIC WORKSTATION

R.R. BABCOCK
R.E. McINERNEY
B.T. BANCROFT
R.J. RAISTRICK

1 September 1989

Approved for public release; distribution unlimited.

AEROSPACE INSTRUMENTATION DIVISION PROJECT 7659
GEOPHYSICS LABORATORY
HANSCOM AFB, MA 01731-5000

"This technical report has been reviewed and is approved for publication"

ROBERT E. McINERNEY
Branch Chief
Data Systems Branch
Aerospace Engineering Division

FOR THE COMMANDER

C. NEALON STARK
Director
Aerospace Engineering Division

This report has been reviewed by the ESD Public Affairs Office (PA) and is releasable to the National Technical Information Service (NTIS).

Qualified requestors may obtain additional copies from the Defense Technical Information Center. All others should apply to the National Technical Information Service.

If your address has changed, or if you wish to be removed from the mailing list, or if the addressee is no longer employed by your organization, please notify GL/IMA, Hanscom AFB, MA 01731. This will assist us in maintaing a current mailing list.

Do not return copies of this report unless contractual obligations or notices on a specific document requires that it be returned.

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188

1. AGENCY USE ONLY (*Leave blank*)	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED
	1 September 1989	Interim Scientific

4. TITLE AND SUBTITLE	5. FUNDING NUMBERS
The Zenith Z-248 as a Scientific Workstation	PE62101F PR76590500
6. AUTHOR(S) R.R. Babcock R.E. McInerney B.T. Bancroft R.J. Raistrick	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Geophysics Laboratory (LCY) Hanscom AFB, MA 01731-5000	8. PERFORMING ORGANIZATION REPORT NUMBER GL-TR-89-0174 IP No. 338
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSORING/MONITORING AGENCY REPORT NUMBER
11. SUPPLEMENTARY NOTES	
12a. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release, distribution unlimited	12b. DISTRIBUTION CODE

13. ABSTRACT *(Maximum 200 words)–(Continue on reverse if necessary and identify by block number)*

This report reviews the development of the desktop computer and workstation for scientific data analysis and develops the basic concepts of connectivity, processing power, high-resolution graphical displays, and large local data storage that support the creation of the scientific workstation approach to geophysical data analysis. The report reviews the capabilities of the basic Z-248 configuration and identifies key enhancements in memory expansion, high resolution displays, added storage, and networking capabilities that can turn the Z-248 into a scientific workstation.

14. SUBJECT TERMS	15. NUMBER OF PAGES
Zenith Z-248 Personal computer Scientific workstation	52
	16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
Unclassified	Unclassified	Unclassified	SAR

NSN 7540-01-280-5500

Standard Form 298 (Rev. 2-89)
Prescribed by ANSI Std. 239-18
298-102

PREFACE

The Data Systems Branch at the Air Force Systems Command's Geophysics Laboratory, Aerospace Engineering Division (GL/LCY), provides data processing and computational data analysis support to the Geophysics Laboratory's scientists and engineers for the interpretation and exploitation of results from the experimental programs that play a critical role in the Laboratory's space science research programs. An integral part of this support has been the creation of the Scientific Satellite Data Analysis System (SSDAS). The SSDAS is dynamic in its approach to meeting the requirements for processing large volumes of scientific measurement data. Advances in the complexity and sophistication of the experimental programs demand the continuous development and implementation of new concepts and new technologies for data processing to meet the needs of the research scientists and the systems development community. The sudden explosion of new computer technology, such as reasonably priced and powerful mini and micro-computers, the ready availability of local and wide area networks for communication, advances in data base management and common data formatting, and the introduction of high-powered graphics workstations, has introduced a whole new era in scientific data analysis and visualization. The Data Systems Branch is implementing a number of major changes in its approach to managing the processing and analysis of large scientific data bases that take advantage of this new technology to improve the efficiency and effectiveness of the SSDAS. This Technical Report is one in a series of Technical Reports and Technical Memoranda produced by the Branch to describe the application of this new technology to meet the Laboratory's data processing needs and to document the new support capbilities and resources that are available.

CONTENTS

1.0 INTRODUCTION

1.1 Purpose of the Report

The rapid development of the desktop computer, plus major technological advances in the design and capabilities of scientific computers, have ushered in a whole new era for scientific data analysis and interpretation. The purpose of this Technical Report is to provide the research scientist at the Geophysics Laboratory (GL) an overview of how this new technology can be exploited through the concepts involved in the utilization of scientific computer workstations. The report will not go into the technical details of computer systems hardware design, but will focus on the basic concepts and capabilities of this new technology. To use the power of this technology effectively, scientists need to understand the concepts of a workstation approach and its advantages for data analysis and research. However, as the technology available for creating a scientific workstation grows more sophisticated and complex, it becomes more and more difficult for the non-expert to know how to choose the right system for his needs. Even for those who try to remain current, selecting between the ever increasing number of choices and options can become a mind-numbing exercise.

This report will explain the scientific workstation concept from the research scientist's point-of-view and describe how the Zenith Z-248 desktop computer can be enhanced to exploit the concepts. The basic needs and requirements of the scientist will be related to the performance capabilities of available hardware and software systems. Once these capabilities are defined, a computer professional can then design the specific hardware and software configuration to meet the researcher's needs.

Received for publication 29 August 1989

This first section of the report includes a brief review of the development of the basic desktop computer in scientific research applications and how it has evolved to be a complete workstation that can be an integral part of the research environment. In Section 2 the concepts and philosophy of the scientific workstation approach to data analysis are explained.

Section 3 develops the workstation concept around a configuration based on the Zenith Z-248 desktop computer systems that are common at the Laboratory. A basic workstation configuration, including the computer and display systems, memory sizes, data storage needs, communication capabilities, and input/output capabilities are described in general terms. The actual design of the hardware and software components for a Z-248 workstation will depend on the individual scientist's needs. However, the Z-248 (or similar system) is a basic configuration that provides the principal capabilities of a workstation. The range of enhancements and special features available for these systems is almost endless. However, there are a few classes of enhancements and options, such as networking with other workstations and computer systems, special imagery display systems, and peripheral storage devices, that are generally applicable to the types of experimental science that makes up a significant part of the Laboratory's research effort. Without overwhelming the reader with technical specifications, these enhancements will be described in enough detail to clearly identify the advantages they offer and define the unique hardware and software that is needed to implement them. When effectively implemented, the Z-248 based workstation can meet many of the demands for scientific analysis and data processing.

1.2 Evolution of the Scientific Workstation

In the 1960's, prior to the advent of the personal computer, most scientific computational work and computer analysis was performed on large, mainframe computers operated at a central site. The scientist submitted data analysis and applications programs, often in the form of punched cards, to computer operators to have the programs run in a "batch" mode. This process was slow. The turnaround was often long, not only because of the demand for use of the systems, but because the early generation computers were slow, single-user systems and their input and output capabilities were limited.

The advent of remote terminals connected to mainframes that had new, multi-user operating systems gave the scientist direct access to the mainframe computer. The multi-user operating systems allowed the computer to support more than one user at a time, introducing the ability to timeshare the mainframe computer and provide more efficient operations. Most computers supported on-line editors, so researchers could write and edit the analysis programs and data bases that were stored on the main computer system. However, the scientist still did not have the capability to interact quickly and easily with the programs and data. The remote or "dumb" terminal was essentially just an automated replacement for the computer operator. Programs were still typically run in a "batch" mode process. The researcher had to put the program in the queue and then later review the output files stored on the mainframe or pore over pages of computer output. Requirements that demanded sophisticated output, such as graphical displays and plots, color displays, and special media output, such as microfiche or 35mm slides, and large data bases, required on-site access to the mainframe computer. Operating systems continue to grow in sophistication. Today most mainframe operating systems, like VMS on the VAX

systems and NOS/VE on the Cyber, support interactive processing of applications programs between a remote terminal user and the computer. However, this mode of operation depends on network access. For dial-in type systems, which have low data transfer rates, interactive processing over these networks can be slow and cumbersome. Higher data rate networks like the Laboratory's Ethernet are fast and efficient, but can be expensive to install, and they can also become saturated. For large organizations with a high demand for interactive data analysis, using the central site mainframe as the primary source of support for this interactive processing is not really an effective use of the power of the mainframe systems.

In the late 1970's and early 1980's the rapid development and acceptance of the desktop computer brought the potential for a major revolution in scientific data analysis. However, the computational power required to meet many of the needs of the scientific community could just not be met by the early PC's with their 8-bit word size, 8-bit input/output bus, and their relatively slow processor speeds. The first IBM PC used an Intel 8088 chip that had a 16-bit processor, but it still communicated in 8-bits. Early PCs had only monochrome displays which could only display alphanumeric characters. A special graphics display processor was required to create the graphical display. The early Computer Graphics Adapter (CGA) processors available on the first IBM PC as an option achieved a four color graphics screen image resolution of about 640 by 200 picture elements, or pixels. (New, high resolution graphics systems that are standard on the latest workstations have 1024 by 1024 pixels, with even higher resolution systems starting to appear on the market.) The very rudimentary display capabilities of the early PC just could not meet the demands of the sophisticated and complex data displays needed for geophysical experimental research programs. Since a great deal of the computer work done by scientists is computationally intensive, the early personal computers,

like the IBM PC and the Apple MacIntosh, found their greatest applications in the business and administrative arenas, where processing requirements focused more on data entry and interaction with data bases. In fact, the commercial success of business application programs such as spreadsheet programs like Supercalc and Lotus 1-2-3, data base management programs like DBase, and several word processing programs, like WORD, Wordstar, and Wordperfect, really provided the impetus for bringing the personal computer out of the hobbyist's workshop and made it an accessible and useful tool for the "non-expert."

The introduction of the IBM AT-type architecture built around the faster, 16-bit Intel 80286 chip with a 16-bit communication bus and the development of inexpensive mass storage hard disks brought a powerful computer to the scientists' desk that is capable of storing and analyzing much of their data in a convenient, readily available form. The Air Force contract for the Zenith Z-248 computer systems, with a large number of options, made access to these systems easier, cheaper, and faster. However, the PC continued to be used primarily for management applications. There were very few sophisticated data analysis and display programs for the scientific community designed to operate on the PC. The Enhanced Graphics Adapter systems that came with the Z-248 increased the pixel resolution of the displays and added more colors, but it still could not compete with systems like the Tektronix graphics display systems that were supported by the mainframe computers. Also, learning to program the PC could often require a major investment of the researcher's time for training and experimentation. These shortcomings meant that the PC had not yet evolved into a really efficient and effective research tool.

Even before the explosion of the desktop computer phenomena, engineering workstations were being developed that exploited the availability of sophisticated, high resolution, interactive graphics for Computer Aided Engineering (CAE) and Computer Aided Design (CAD) problems. These workstations were essentially interactive remote terminals with very high resolution display systems networked to a smaller mainframe computer, like a VAX 11/780 or a Prime minicomputer. Although quite powerful, they suffered from the same problems as any remote terminal system -- as the network became loaded down, the performance decreased rapidly. In the early 1980's, Apollo Corp and Sun began developing stand-alone workstations built around 16-bit processors like the Motorola 68000 series. These systems incorporated high-resolution graphics display systems and multi-user/multi-tasking operating systems to provide the power of a mainframe computer in a package that could easily fit on a desk top and with a price range that was inexpensive by mainframe standards.

The rapid advances in technology and the competition for the expanding workstation market have quickly brought a convergence in the capabilities and price of desktop personal computers and the micro-computer based workstations. The next generation workstation that is being marketed by Digital Equipment Corporation (DEC), Hewlett-Packard, Sun Microsystems, Silicon Graphics, Inc., and other manufacturers fully implements these workstation concepts with the latest state-of-the-science micro-computer technology. Figure 1 illustrates the evolution and convergence of these systems.

Scientific applications at GL are often heavily involved in experimental data analysis and interpretation. There are important aspects to designing a "workstation" for supporting experimental research that separates the functionality of a workstation

from the stand-alone personal computer. Before discussing how to turn a Z-248 into a fully capable workstation, an important first step in creating the final configuration of the system is the definition of the concept and the characteristics of a workstation.

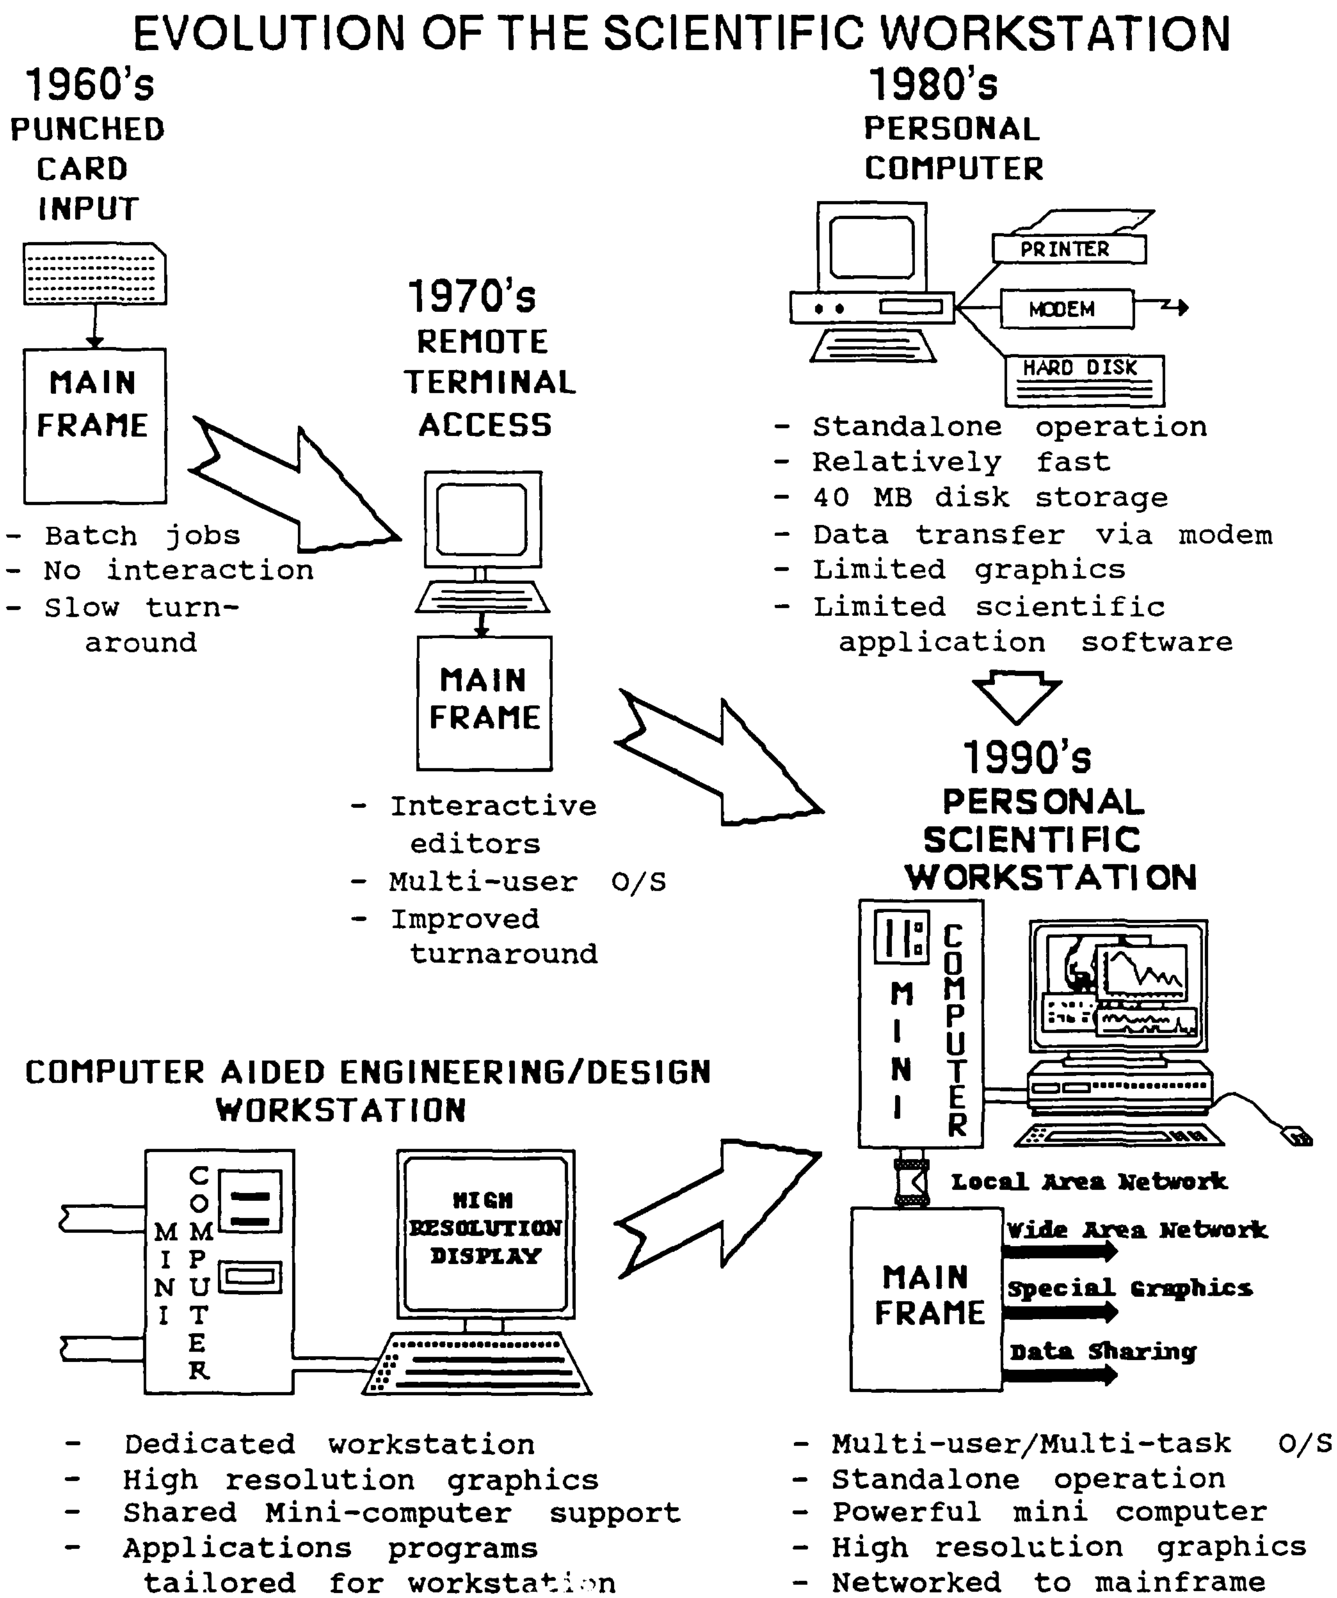

Figure 1. Evolution of the Scientific Workstation

2. THE SCIENTIFIC WORKSTATION APPROACH

2.1 The Workstation Concept

Before deciding if a workstation approach to scientific data analysis is right for any particular program, the concepts involved in a workstation need to be defined. What do we mean by a "workstation," and how are the capabilities of the workstation different from the capabilities of a stand-alone personal computer or a remote terminal linked to a mainframe computer? There are four principal characteristics of a workstation that separate it from the other two working environments:

1) Processing Power - the workstation has a powerful computer processor to provide stand-alone data computation and manipulation capabilities and a large random access memory to support large programs and rapid display of data.

2) High Resolution Graphics - the workstation supports a high level graphical interface with the user to provide rapid, high resolution, multi-dimensional visualization of data for display, analysis, and interpretation.

3) Data Storage - the workstation includes sufficient random access storage capability to allow fast, efficient interaction with working data bases.

4) Connectivity - the workstation is networked to other computers to share resources, including hardware, software, and data.

2.1.1 PROCESSOR POWER

The computer chip is the engine that drives the computer and determines the overall capabilities of the workstation system. The processor provides the computational power that supports the stand-alone data processing objectives of the workstation. The processor also establishes the baseline for adding enhanced capabilities, such as high resolution graphics, multi-user operating systems, and large memory sizes. In the actual acquisition of a workstation, the specific computer chip used is essentially irrelevant to the researcher. The real interest is in performance requirements. The older technology used in chips like the Intel 80286 can still support many of the requirements for connectivity and high resolution displays when augmented by special adapters, add-in cards, or coprocessors. For pure data processing, relatively straight forward computational tasks, and programs with limited data transfer and graphical display requirements, these older systems are probably sufficient. The commercially available workstations use faster, more powerful chips, typically include coprocessors for hardwired floating point calculations, larger standard memory boards, and the chips are designed specifically to support multi-user and multi-tasking operating systems. This next generation processor is essential for more powerful applications.

2.1.2 HIGH RESOLUTION GRAPHICS

High resolution graphics is the key to making the workstation a powerful diagnostic and analysis tool. The developers of CAD and CAE systems quickly discovered that low resolution displays just could not clearly represent the intricate graphics of complicated engineering systems. High resolution graphics were essential for

effective visualization of these complex systems. Also, the power of high resolution, three-dimensional (3D) displays greatly increased the value of the graphics workstation as a design tool. Effective implementation of high resolution graphics requires sophisticated graphics processors and dedicated graphics drivers to get the fine detail and rapid display of these high performance systems. The latest high resolution graphics systems with 1024 by 1024 pixel resolution allow precise representation of the detail of complex designs and also make possible the high fidelity display of multiple images on the screen. This improved resolution has two impacts on the use of workstations. First, many of the new imaging systems being designed and flown on space experiments can produce imagery with resolution as high as 1024 by 1024 pixels. To display a full image, high resolution is essential. Second, the display of multiple graphics and imagery on a single screen makes the workstation an efficient analysis tool. This multiple display approach is typically accomplished using a windowing technique. Each display, image, or graph appears on the screen in its own "window." The researcher can display several pieces of analyzed data on one screen for easy comparison. Windows can be quickly opened and closed, temporarily hidden from view, or reduced in size and moved about the screen to give the analyst total control in creating the "picture" that best supports interpretation of the data. This multiple display visualization technique can be accomplished effectively only with the newer, high resolution systems. Older technology, such as available on the Z-248, rapidly loses fidelity as the number of windows increases because of the low pixel resolution.

2.1.3 DATA STORAGE

The availability of inexpensive internal hard disks and fast access peripheral storage devices, like optical disks, allow for the storage of large working data bases and the

creation of large programs at the scientist's workstation. This means the researcher can have the data and programs readily at hand for use without the constraints of network access availability or depending on central site operator support to load magnetic tapes or disks. The central computer facility still plays an important role for the archival of large data bases and providing access to multiple data bases for many different users. But for active visualization and interpretation of working data bases, the expanded capabilities of the workstation for data storage is one more factor that adds to the efficiency and effectiveness of the workstation concept for scientific data analysis.

2.1.4 CONNECTIVITY

At the most basic level, connectivity means communication. But in application, it means access and sharing of resources and the synergism that this sharing supports. Connectivity is an important characteristic for making the complete workstation concept successful, particularly when used for large, multi-discipline experimental programs. Connectivity is typically accomplished through networking, and there are several levels of networking that can be implemented in the workstation configuration.

Local Area Network (LAN) - This is a communication network within an organization. A LAN provides two important capabilities. The first capability is access to a mainframe for support. GL has a high capacity Ethernet network that provides an easy, direct connection to the mainframe computers at the central site. Access to the network services of a mainframe computer is a critical consideration in the design of a workstation configuration. Many experimental programs

have large databases, but researchers do not require continuous, daily access to the entire data base. The central site mainframe acts as a file server. The data can be stored or archived at a central computer facility where large, expensive peripheral storage devices can be more efficiently and effectively utilized. The researcher easily and quickly retrieves the limited working data bases needed for any particular problem through the LAN, thus eliminating the need for excessive, dedicated storage capability with the workstation. The central site also has a regular back-up procedure for all data stored on the disk. This support eliminates the need for developing back-up procedures and hardware for large data bases stored with the local workstation, or what can be worse, losing an entire data base because of operator errors or hardware failure. The mainframe can also be a computation server providing the workstation access as a remote terminal connected through the LAN for the use of specialized applications or analysis programs that are only available on larger computers, for use of specialized display or output systems, or for running computationally complex problems that are more efficient on a mainframe. The second capability a LAN provides is direct communication with any other station on the network including rapid peer-to-peer communication within the Laboratory for data transfer and data sharing.

Work Group Networks - A workstation can also be a network server for a smaller work group network. The primary workstation becomes a centralized resource for data storage and analysis programs, while the add-in terminals and workstations provide added analysis

productivity by expanding the interaction and communication within a working group focused on one research program.

Wide Area Networks - Wide area networks, such as the Defense Data Network (DDN) and NASA's Space Physics Analysis Network (SPAN) provide connectivity between GL and other research organizations and facilities. Besides providing peer-to-peer communication for data transfer and data sharing, this network also provides access to the data processing, analysis, and computational capabilities of these other activities. At GL the central site acts as a communication server and is the gateway to these networks. Any workstation connected to the central site through the LAN has access to these networks. Special networking arrangements can also be created using dedicated communications lines, such as the data line to the Cray at the Weapons Laboratory.

2.2 Exploiting the Scientific Workstation

Figure 2 summarizes the basic characteristics and capabilities of the next generation workstation concept. The workstation approach to data processing and analysis brings a powerful computer to the desk of the individual scientist, while providing fast access to the capabilities of a main frame computer and maximizing efficiency through the sharing of resources and data. Connectivity supports and encourages the synergism and cohesiveness of a small work group environment, while at the same time providing unlimited access for data sharing through peer-to-peer communications. In today's research arena with large experimental programs often involving science working groups with team members hundreds of miles apart, the

need for rapid data transfer and the quick dissemination of information becomes critical. The workstation, when properly connected to available local and wide area networks, allows easy and fast peer-to-peer communication. The individual scientist has the power to meet individual analysis needs in a single-user environment, while still being able to quickly share data and results with colleagues and associates.

Because the cost of the state-of-the-science hardware is relatively small, the workstation can be economically designed to support a single work effort, a research project, or an entire experimental program. The workstation configuration, including the primary computer, storage devices, network design, locally connected terminals, display systems, operating systems, and applications software can all be designed up-front to meet the specific needs of the project. With proper planning, the data format structure and unique analysis programs can be designed and developed before the experiments are performed, so the workstation is ready to support the data reduction and analysis as soon as the experimental data is available.

Using a workstation approach to experimental data analysis has several advantages. First, the connectivity of the workstation allows the scientist almost unlimited access to any data needed to support the research. If the data base from the experiment is large, it can be stored at a central computer facility, such as the central site computer facility at GL that includes a DEC VAX cluster and Control Data Corporation Cyber mainframe. Selected portions of the data can be extracted from the entire data base for analysis. If data from other experiments or other agencies is required, data sharing can often be accomplished easily and quickly using an existing local area network or wide area networks like the DDN or SPAN.

Once the data is resident in the workstation, the scientist has the computational tools to personally perform the analysis. Modern workstations come with highly sophisticated operating systems that provide a graphical, user-friendly interface. Even the non-expert can quickly learn how to manipulate the data and perform fairly sophisticated analysis. There are numerous commercially available analysis packages, plus several access and display routines developed for GL-specific systems and programs, that allow the scientist to quickly and efficiently review the data and select important or significant events for more detailed analysis. This eliminates the drudgery of reviewing stacks of computer printouts, microfiche, film, or plotter output. The scientist's desk no longer needs to be a clutter of charts, graphs, and computer output. The graphical or tabular output is reviewed in "softcopy" on the workstation screen. Significant events or regions of interest can be quickly identified, stored, and easily recalled later for review or, if necessary, printed out for more detailed analysis. The multi-tasking capabilities of the workstation operating systems allows several displays, images or graphs of the data to appear on the screen to speed the analysis. Of course, high resolution graphics screens and hardware graphics processors and drivers are needed to fully utilize this multi-display technique effectively.

The workstation approach to data analysis means researchers have ready access to the data and the analysis tools at their convenience. They are not restricted by the schedules and conflicts of other computing centers; networking gives them access to the power of a mainframe for the really complex and computationally intensive problems that mainframe computers are most efficient at solving. Most of the commercially available workstations come with the highly developed and sophisticated graphical user interface programs with multiple windows and detailed menus that make the use of the systems a less demanding and complicated task.

With a small investment in learning the basic start-up procedures and the implementation of any of the newly-available applications programs designed specifically for science data analysis, a workstation can be created that can be exploited by all scientists, no matter what the level of their computer literacy.

SCIENTIFIC WORKSTATION CHARACTERISTICS

HIGH RESOLUTION GRAPHICS

HIGH RESOLUTION IMAGERY
GRAPHICAL USER INTERFACE
MULTIPLE WINDOW DISPLAY FOR VISUALIZATION OF DATA
PSEUDOCOLOR ANALYSIS

PROCESSING POWER

MULTI-USER/MULTI-TASKING OPERATING SYSTEMS
COMPUTATIONALLY FAST
EXPANDED MEMORY FOR LARGER APPLICATIONS
SUPPORTS LARGER DATA BASES
INTERFACE WITH MORE PERIPHERALS

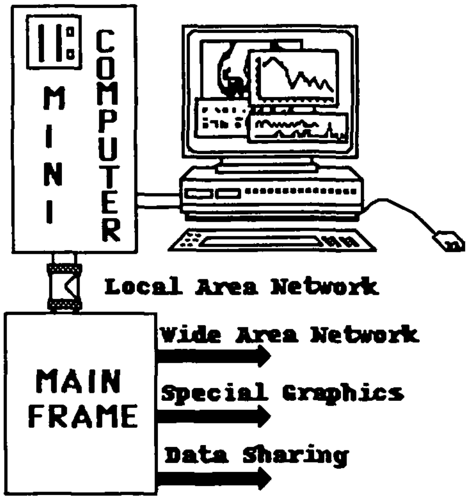

DATA STORAGE

FAST, RANDOM ACCESS STORAGE
WORKING DATA BASE AVAILABILITY
RAPID ACCESS FOR VISUALIZATION

CONNECTIVITY

LOCAL AREA NETWORK
- MAIN FRAME ACCESS
- SHARED RESOURCES
- SPECIAL OUTPUT DEVICES
- DATA STORAGE
- DATA TRANSFER/SHARING
- APPLICATIONS PROGRAMS ACCESS
- PEER-TO-PEER COMMUNICATION

WIDE AREA NETWORK
- COMMUNICATION & DATA SHARING
- EXPAND FACILITIES SUPPORT
- ACCESS OTHER APPLICATIONS

WORK GROUP NETWORK
- INCREASE LOCAL INTERACTION AND PRODUCTIVITY
- EFFICIENT USE OF RESOURCES

Figure 2. Summary of Scientific Workstation Characteristics

3. THE ZENITH Z-248 BASED WORKSTATION

3.1 Performance Characteristics

Before discussing the specifics of enhancing the Z-248 desktop computer to make it a workstation, we need to first categorize the basic performance capabilities expected of any workstation or personal computer system and relate these to the hardware and software components that make up the computer. In later sections, several enhancements for the Z-248 will be discussed, and it is important to understand what performance capabilities or deficiencies these enhancements affect. More technical details on some of the concepts and components are given in the Glossary.

The engine that drives the workstation is the computer chip. The chip consists of two principal components; the central processing unit (CPU) does all the computational work and controls the computer's operation, and the main memory stores data. The speed and performance characteristics of the computer depend on such design characteristics as the clock speed of the processor, size of the words used (number of bits in a word), memory access time, the sophistication and complexity of the instructions used by the processor, the presence of an add-in co-processor to do floating point calculations, the size of the words of data that can be passed through the bus that connects the computer to the storage devices (for example a hard disk), and the amount of memory available. (See the Appendix for more details on factors that affect speed and performance.) Computer chip designers make compromises among these various options in developing a class or family of chips. The Intel 80x86 series Computer chips are used in PC class systems. (The Intel 8088 was in the original IBM PC, the Intel 80286 is in the IBM-AT and

the Z-248, and the newest PC's use the Intel 80386 chip). The Motorola 68000 series chips are used in MacIntosh systems.

The next major component that is critical to the performance of the complete computer system is the software operating system. The operating system is the program that acts as the interface between the user and the processor. The operating system handles the loading of programs, interacts with peripheral devices, looks for interrupts from the keyboard and other devices, schedules the use of memory and maintains a memory map of where things are stored, and acts as a general traffic cop for all the processes that are going on in the computer. No matter how powerful the computer chip is, the computer system can only perform to the level that the operating system allows. The Z-248 and most IBM PC-type systems use the Disk Operating System (DOS) designed by Microsoft Corporation specifically for IBM and the Intel 8088 chip that was in the original IBM PC. DOS is a single user, single task operating system. Many of its shortcomings in today's computing environment, such as the limited memory that can be used for programs (640 kilobytes) and the size of a storage drive (32 megabytes) are a legacy of original design of DOS and the structure of the Intel 8088 chip. There are other operating systems that can be used on the Z-248, including the Unix systems and IBM's OS/2, both of which are multi-tasking systems, and several commercially available extensions of DOS that solve some of the shortfalls of DOS.

There are several important peripheral devices that are critical to the overall performance capabilities of the Z-248 system. Random access storage devices, such as floppy disks and hard disks, augment memory for storage of data and programs. Hard disks provide high volume and fairly rapid access storage. The floppy disk drives are slower with less capacity, but are ideal for transport and installation of

programs and data. The video display is critical to the workstation concept for rapid, flexible display of softcopy products used in data analysis. The major factors to consider in video display systems are pixel resolution, color capability, and the capabilities of the graphics interface card that drives the video display. A high resolution screen is wasted if the graphics driver cannot utilize (drive) all the pixels. The Enhanced Graphics Adapter (EGA) graphics card found in the standard Z-248 configuration can support a nominal 640 x 350 screen image and display sixteen colors at a time. High resolution workstations with the latest display systems have 1024 x 1024 pixel resolution and can display several million colors.

A modem is an important communications option that allows the computer to transmit data over telephone lines. At GL, a computer with a modem can dial into the central site and get access to the broadband network. With this access the user can utilize the VAX cluster for data storage, communicate with other scientists in GL, obtain access to DDN and other wide area networks, and utilize the other capabilities of the central site system.

The printer provides a hard copy output capability. Most Z-248 systems have an ALPS P2000G Desktop II dot matrix printers that performs both alphanumeric and graphics printing. The Z-248 can support newer laser printers that provide higher quality print and enhanced graphics capabilities. Many of the Z-248 systems in the Laboratory already have Hewlett-Packard Laserjet or DEC LN03 laser printers

3.2 Z-248 System Configuration

Table 1 lists the basic hardware components that make up a typical Z-248 system. How do these hardware components translate into workstation capabilities and what

level of performance can be achieved with this basic system? Let's look at each of the major functions that make up the workstation concept of data analysis and assess the performance.

Processor Power: The Intel 80286 processor is reasonably fast and capable of handling complex computational problems. The addition of the math coprocessor for floating point calculations is essential for most scientific calculations. The 640 kilobytes of random access memory are essentially the maximum that the DOS operating can utilize. When DOS is the resident operating system, the Z-248 cannot support multiple-user operations or multiple tasking. Also, DOS is not particularly

Table 1 Typical Z-248 Configuration

Component	Type
CPU	80286
Coprocessor	80287
Clock speed	8 MHz
Instruction Speed (MIPS)	0.5-1.0
Memory	640 Kbytes
Hard Disk	40 MBytes
Floppy Disk	1-1.2 MByte
Networking	Modem
Operating System	DOS
Graphics	EGA/Color 14" Monitor
I/O	AT bus

user friendly for the non-expert, so a higher level, graphical user interface program, such as commercially available programs like Microsoft Windows 286, DESQview or the Norton Commander are an important option to consider. Multi-tasking can be achieved on the 80286 chip, but it takes a special software operating system or application program that handles the scheduling of the multiple tasks and manages the use of memory. The 16-bit word size is adequate for most applications and data processing. However, for problems that require double precision calculations, that have large data bases that must be accessed continuously, or that use analysis or modeling programs that are computationally intensive, the 16-bit word size becomes a significant factor in degrading accuracy and speed. The use of double precision arithmetic to increase accuracy can add a significant penalty in performance.

High Resolution Graphics: Most of the Z-248 systems have an EGA color display that has a nominal 640 by 350 pixel screen image resolution. For single image displays, this resolution is normally quite sufficient. The screen can easily support multiple window displays. However, if more than two windows are displayed, the fidelity of the display in each window decreases rapidly. If the windows contain alphanumeric or tabular data, the impact is not significant. But for data intensive graphical displays, like multi-spectral analysis in pseudocolors, the display can quickly become unusable if the window is reduced in size for the simultaneous display of multiple windows.

Data Storage: The 360 kilobyte floppy disk drive and either 20 or 40 megabyte internal hard disk are a relatively standard configuration that provides basic desktop processing capabilities. However, for programs with large experimental data bases and a large number of analysis and applications programs, 40 megabytes of storage can be quickly filled.

Connectivity: Connectivity and communication within the Laboratory are typically achieved through the Broadband asynchronous channel hosted by the VAX cluster that uses the multi-channel cable plant. This link provides access to the central site mainframe for data storage, data transfer, and data sharing. The LAN also provides nodes for wide area network access to the DDN and SPAN through the central site node. A modem can also be used to dial into other computer systems and research facilities to further expand wide area networking. The problem with these links is that the transmission rate is slow (typically 9600 Baud for the LAN and 1200 or 2400 Baud for a modem). A modem connection also is limited by the number of connections a site can support, the availability of the system, and the fact that it ties up a phone line.

3.2 Enhancements for the Z-248

There are several specific enhancements to the basic Z-248 system that can significantly increase workstation capabilities of the system in processing power, high resolution graphics, data storage, and connectivity. Examples of these enhancements are summarized in Table 2. Most of these enhancements can be obtained for a relatively modest investment. However, the cost of a next generation workstation is almost comparable to the cost of enhancing an older technology Z-248 to full workstation capability. With more powerful processors, high resolution display systems, and specially designed graphical user interfaces that make them much easier to use, the new workstations not only have an exceptional price/performance ratio, but the competitive pricing is making them increasingly more affordable.

The next few sections will discuss examples of some of the enhancements to a Z-248 that will bring it up to the workstation configuration as defined in Section 2. A complete chapter could be dedicated to describing the options and variations for each of these enhancements. The detailed discussion of each type of enhancement can be found in excellent, in depth review articles in any of the computer magazines, and that detail will not be reproduced here. This section will only summarize the key factors to be considered in choosing an enhancement and provide a general estimate of the costs. Issues that should be considered when comparing the trade-off between enhancing a Z-248 and upgrading to a next generation workstation is discussed in the last section.

Table 2. Enhancements for the Z-248 Workstation

Enhancement	Improved Capability	Deficiency Satisfied
Connectivity *Ethernet Card to connect to the GL LAN*	Increase speed of data *transfer and communi-*cation	- Faster access to large data bases stored on the central site or at other facilities - Better utilization of central site storage
Processing Power Expanded Memory	Increased memory for larger programs	- Run larger applications programs
Data Storage Additional Hard Disk	Store larger data bases and more programs	- Reduce time to access data from central site
High Resolution Graphics High Resolution Monitors	Higher fidelity in displays More effective use of multiple window displays	- Fidelity and resolution of displays - Inefficiency in amount of data displayed
Graphical Interface/ Presentation Software	User friendly operating environment	- Reduce complexity of operating the computer

3.2.1 CONNECTIVITY ENHANCEMENT - ETHERNET INTERFACE ADD-IN CARDS

The most important enhancement to the Z-248 system that brings more effective and efficient workstation capabilities is an add-in card for interfacing to the Ethernet communication network already available in GL. The coaxial cable used for Ethernet baseband network allows a nominal 10 megabits per second bandwidth for transmission of data versus the 1200 to 9600 bits per second capability of the telephone lines used for a dial-up modem or the broadband system. Because of the overhead involved in the transmission protocols required for a multi-user system and the natural limits of the Z-248 communication bus, the actual efficiency of Ethernet is reduced to a few hundred kilobits per second. But this is still significantly faster than the broadband asynchronous channel of the LAN. For applications with large data bases stored at the central site main frame, for applications that utilize frequent data transmissions for data sharing, or for research programs that require frequent access to the main frame computer for applications programs that cannot operate on the Z-248, the Ethernet connection can greatly increase performance and usefulness of a Z-248 workstation.

A complete Ethernet LAN provides the cabling for the data transmission, the basic protocols that are used for transmitting the data, and a network server, like the central site VAX cluster. An Ethernet LAN can be established at any level that a server is available, such as a Division or Branch. In order to implement an Ethernet connection to the central site VAX cluster, two components are required. First is the hardware Ethernet card that is added to the computer. This card provides for the connection to the network and contains a set of processor chips that implement the communication protocols of the Ethernet system. Some of the cards available

commercially include storage buffers to hold data for transmission and their own stand-alone processors to accomplish the data handling and relieve the computer processor of this task.

The second component required is the software program to manage the data transmission process and interface with the workstation user. There are a number of commercially available network software systems that implement an Ethernet LAN, and many common bundled with the add-in card and the software. DEC has its own proprietary network architecture system called DECnet for VAX systems that implements the Ethernet capabilities. DEC also developed a special software package called DECnet Personal Computing Systems Architecture (PCSA) network software that is designed specifically to support networking a personal computer like the Z-248 with both the DEC Ethernet and asynchronous connections. The DEC/PCSA Client network software essentially allows the Z-248 to use a VAX host server computer like another disk drive for data and program storage, besides providing the other connectivity capabilities that come with any LAN.

There are commercially available networking systems that allow establishing a work group network within an office or branch using just the PC's. These systems usually include the add-in cards and proprietary software for managing and operating the system. Some of these systems use configurations other than the Ethernet topology, such as token ring and star topologies (See the Glossary for more details.) (Because DOS does not support multi-user and multi-tasking capabilities, these PC networks tend to tax the performance of the host, or server, and degrade the overall efficiency of the network.)

The retail cost of the add-in Ethernet cards range from $295 for a basic capability up to $525 for faster cards with extra memory storage on the card. The license for the DEC/PCSA Client software is $250. (The basic software is already available in the Laboratory.) A combination DEC DEPCA Ethernet card and the DEC/PCSA Client software license has a current (June 1989) retail price of $495.

3.2.2 DATA STORAGE - ADDITIONAL RANDOM ACCESS HARD DISKS

Adding data storage is the easiest enhancement to the Z-248 workstation. Additional local storage reduces the requirement to access data from the central site mainframe, which is particularly slow if an Ethernet connection is not available. The additional storage also allows for more programs and data to be stored locally for analysis and processing.

Most Z-248 systems come with either a 20 or 40 megabyte hard disk. Despite the initial assumption that this is far more storage than could ever be required, the disk is usually quickly filled, particularly if there are several users of the system. The addition of a second hard disk is easy and fairly inexpensive. Additional hard disks with capacities from 20 to 330 megabytes are available with access times from 18 to 65 milliseconds. Access time is important if the applications programs require frequent access to the disk for data retrieval and storage.
Besides the hard disk itself, the computer may also need a controller. This is the device in the computer that manages the data flow and interacts through the communication bus with the processor and operating system. Most Z-248 systems have only one hard disk controller, so the additional hard disk system must include a controller. One way to solve this problem is to get what is known as a hard card which includes the disk drive and controller in one add-in card. The additional hard

disk can either be installed inside the computer or be an external add-on. Prices range from about $300 for a 40 megabyte, 40 millisecond access time system, to over $1000 for large capacity, high speed systems that use the faster Small Computer Systems Interface (SCSI) adapters.

Another alternative to a hard disk is an optical disk drive. These systems have the advantage of removable media, like a floppy disk, so they won't be saturated like a hard disk, and they have the fast access of a hard disk. Installation requires a host adapter add-in card plus the external disk drive. Compact Disk - Read Only Memory (CD-ROM) can only read disks created by another system and are really only useful for reading archived data. These systems cost about $1000. Write Once - Read Many (WORM) systems that allow creation of stored data for permanent storage (since they cannot be erased) cost from $5000 to $7000 for the host adapter controller card and the drive. The cost depends on the amount of data that can be stored and the speed of the host adapter in transferring the data. WORMs can hold from 800 megabytes up to 1.2 gigabytes of data, which is equivalent to over 20 40-megabyte hard disks. Because this technology is fairly new and the standards are not well established, there can sometimes be difficulties in properly connecting the systems. Unlike floppy disks the optical disks are not always easily interchangeable between different vendor's drives . Advance planning and a little homework is essential to be sure the right system is acquired and properly installed.

3.2.3 PROCESSING POWER - EXTENDED MEMORY

The basic DOS design limits the size of memory that DOS can manage to 1 megabyte. The first 640 kilobytes of memory is reserved for DOS and applications programs, and the last 384 kilobytes, referred to as expanded memory, is nominally

used for video memory and DOS overhead functions. This limit is an artifact of the original 20-bit address size of the Intel 8088 processor that the early IBM PC's used for address storage. The newer Intel 80286 processor in the Z-248 has the capability to address up to 16 Megabytes of memory. The addition of memory above the 1 megabyte limit is known as extended memory. This increased memory can be utilized to store more data in memory in a RAM-disk or cache to reduce the requirement for disk access to read data. The result is faster running programs. The added memory can also be used for spooling print output that allow large printing jobs to be essentially done off-line. With the use of spooling, the processor is not held up with I/O processing, but can return to data processing and computation. Extended memory also permits more effective use of the virtual memory capabilities of the 80286 chip, so larger applications programs can be installed and multi-tasking can also be achieved with an upgraded operating system, but this requires special drivers.

Implementation of extended memory involves two factors; installation of a memory board with more random access memory chips and extended memory manager hardware or a software program that manages the extended memory and interacts with the applications programs. Expanded Memory Specification (EMS) add-in boards provide a hardware driver to make extended memory appear to be expanded memory that DOS can access. The boards with EMS drivers provide the fastest means of utilizing the extended memory, but they also are fairly expensive. Software drivers for EMS are less expensive, but also less efficient. In comparing the add-in boards, price alone should not be the only factor. Comparisons should include the options available for memory management. Table 3 gives a limited list of extended memory boards that are available to indicate the typical expansion sizes available and typical retail costs. The suggested retail price is based on the smallest

memory size increment available. For boards with a minimum of 0 kilobits, the price is for the add-in board with the communications bus alone, and the memory chips are extra.

Table 3 Add-in Extended Memory Boards for the Z-248

Company	Board	Memory Available	Suggested Retail Price
AST Research	RAMpage 286	512k-2MB	$695
Everex	RAM 10000	0-10 MB	$399
Intel	Above Board 286	512K-2MB	$645
Micron Tech	1Mb EMS Board	1-4MB	$895

An important aspect of acuiring extended memory is understanding the implementation of the extended memory manager and the effective utilization of this added memory either through an EMS board or a software system. Full exploitation of extended memory in the Z-248 configuration requires a fairly sophisticated knowledge of the operating principles of the DOS and the PC itself. For the scientist who does not have the time or interest to get involved in these details, the key to successfully exploiting the powers of a Z-248 workstation with extended memory is acquiring a more sophisticated operating system and applications programs that provide a congenial interface with the user and the hidden power for managing the enhancements without operator interaction.

3.2.4 HIGH RESOLUTION GRAPHICS - HIGH RESOLUTION MONITORS

One of the most important enhancements required for creating a really efficient and productive workstation for high quality visualization of complex and multi-sensor

data bases is a large screen, high resolution display system. This requires acquisition of a high resolution color monitor and the graphics adapter card to implement the higher resolution display. The options for high resolution monitors is almost endless. A complete review of these options could be a complete separate section. The best solution for the Z-248 workstation that provides the highest available resolution and is still compatible with most applications programs is the new Super Video Graphics Array (VGA) systems. They provide the 800 by 600 pixel resolution of VGA plus the ability to display up to 16 colors at a time from a palette of 256. The number of colors that can be displayed simultaneously is important for visualization techniques that use multiple windows or for display of multiple data sets. The cost of the graphics cards needed to implement the Super VGA vary depending on the amount of memory available on the add-in card, the downward compatibility capabilities, and other more technical options. Almost all of the cards are downward compatible to be able to support applications programs that utilize CGA, EGA, and regular VGA and Hercules Graphics standards.

The 14" VGA monitors range in cost from $595 for a basic super VGA monitor with few options, up to $1399 for the high end model super VGA monitors with special features that provide automatic sizing between standards (such as EGA to VGA), greater contrast, intensity, and fidelity in the display. Larger format 19" monitors cost up to $3000 or more. The Super VGA graphics add-in adapter cards vary in price from $395 to $695 depending on the amount of video storage memory and other options.

A very important factor to consider is that simply adding a high resolution monitor and graphics card does not automatically provide high resolution displays. It is the applications programs and the graphics drivers in these programs that implement

the hardware capabilities of hardware. This means new programs will have to be written to fully exploit the Super VGA resolution. Many of the latest commercial software packages provide this support, but the DOS application programs developed by the individual scientist will have to be revised to utilize this new capability.

IBM has introduced an 8514/A multi-scanning, interlaced , high resolution monitor system that is compatible with the Z-248 and provides 1024 by 768 resolution displays. The problem with going with this enhancement is that very few software drivers available can implement these capabilities. Acquiring an 8504A system will also require an investment in developing and implementing the software to use the system in any applications programs

3.2 5 GRAPHICAL INTERFACE/PRESENTATION SOFTWARE

The last major enhancement that should be considered for the Z-248 workstation is a graphical user interface program or graphical presentation program. These programs, or shells, provide an additional level of interface between the user and DOS. The latest versions typically include a window environment with menus for each task that give the user options to perform and list available applications programs. This makes the utilization of the system much more user friendly and typically requires less knowledge of the inner workings of the computer and the operating system itself. From an ergonomic point of view, research has shown that a graphical display interface is easier to work with and provides more efficient interaction, with the end result being increased productivity. These programs can also handle many of the interfacing actions required for the other enhancements,

such as the LAN and extended memory, for the operator without direct involvement.

Some of the latest commercially available programs, such Microsoft Windows 286, also provide a capability for implementing an equivalent of a multi-tasking operating systems. The Intel 80286 processor in the Z-248 has the capability to support the virtual memory management that is needed for multi-tasking. It is just the limitations of DOS that keep the Z-248 from being multi-tasking. These special user interface programs allow multiple sessions of DOS to operate within memory and the scheduling of various tasks for each session. Windows 286, DEQView/286, and similar user interface programs cost about $100. Operating system programs that allow multi-tasking, like a UNIX operating system or IBM's OS/2 will operate on a Z-248, but they require their own specially designed applications programs. The price range for these operating systems is $400 to $500, depending on their capabilities and options. Also, OS/2 requires at least 2 megabytes of memory, and really needs 4 megabytes or more to be effective.

3.3 Enhancements or Upgrade - What's the Best Strategy?

The preceding section discussed five of the more cost effective enhancements to the Z-248 system needed to make it an effective workstation that is also efficient and user friendly. There are many other enhancements that can be considered, such as 80386 expansion boards that essentially turn the Z-248 into multi-tasking, 32 bit, 16 MHz, state-of-the-science workstation for about $1900. Each of these enhancements brings an important increase in capability to the basic Z-248 hardware configuration.

The drawback in using the enhancement approach to create a workstation capability is that it often requires a significant level of small computer literacy to choose the compatible systems and install them properly. Besides acquiring the hardware, special software is usually required to effectively utilize the enhanced capabilities. Installation of the software often demands a certain amount of sophistication in programming and operating small computers.

Acquiring more powerful machines, whether by enhancing an existing Z-248 or upgrading to the next generation of workstations, is only part of the solution. The user's lack of sophistication and computer literacy is a major stumbling block in achieving increased productivity with all this new technology. Many scientists do not have the time to learn the details of these systems needed to make an informed decision about what upgrades to acquire or how to effectively implement them once the new hardware and software is installed. An important factor in the decision process is the investment in time and effort that will be required to learn how to use these new capabilities.

Another important issue in the enhancement-versus-upgrade decision is the how to utilize the investment already made in applications software. If significant resources have already been invested in applications programs for the Z-248 that work in the DOS environment, an upgrade might make this investment worthless if the applications programs can not be ported easily to the new system. And what about the existing Z-248 hardware and peripherals? Can the Z-248 system still be connected to a next generation workstation, or will this investment have been wasted?

If the inherent capabilities of the Z-248 system do not meet the scientists data analysis needs, rather than struggle with piecemeal upgrades, a new class of workstation may be a better solution. The decision matrix in Table 4 compares a sampling of the issues and considerations involved in defining various

Table 4 Decision Matrix for Designing a Scientific Workstation

Requirement	Z-248 Workstation Application	Next Generation Workstation Application
Connectivity	-Small group or single user -Limited networking ♦E-mail ♦Submit batch jobs ♦File storage at central site	-Large work groups at many different locations -Heavy demand for network ♦Peer-to-peer communication ♦Wide area network access ♦Extensive data sharing -Local work group network required
Processor Power	-Off-the-shelf DOS applications -Small,locally developed FORTRAN programs -Basic modelling and data analysis	-Multi-tasking/multi-user operating environment -Large, complex analysis & display applications -Computationally complex modelling -Fast, interactive data processing and display -Heavy demands on input/output capabilities -Support for local work group network
Data Storage	-Single experiment -Moderate size data base (<1 megabyte)	-Multi-sensor experiments -Large data bases, large imagery data sources -Multi-dimensional models
Graphical Display	-Mostly tabular displays -Basic X-Y plots used -Limited to two or three displays on a screen -Limited use of color	-High resolution imagery -Multiple display of image, X-Y plot and text data -Pseudocolor important for analysis & interpretation -Interactive graphical display

analysis and visualization tasks that are appropriate for an upgraded Z-248 workstation and a next generation workstation.

The Z-248 workstation is more than capable for many of the experimental and theoretical modelling applications at GL. But for the more complex, multi-sensor, multi-discipline experimental programs of the future, like CIRRIS, CRRES, IBSS, and MSX, the next generation workstation may be an essential component of the total program needed to support the productivity and inter-disciplinary communication that will be critical to the success of these programs. Section 5.0 of this report report discusses the next generation workstation and defines the basic capabilities that should be considered in designing and acquiring one. But before discussing these workstations, this report will first give an overview of the applications software already available at GL for display and visualization of scientific data.

REFERENCES

Baran, Nick (1989) Two Worlds Converge, *Byte*, Vol 14, No 2, p 229-234.

Davis, William S. (1987) *Operating Systems - A Systematic Review*, Addison- Wesley, Reading, MA, pp.539.

Defler, Frank J., Jr. (1989) CONNECTIVITY - Building Workgroup Solutions: Ethernet Cards, *PC Magazine*, Vol 8, No 2, p 155-159.

Kindel, William (1989) How Fast is Fast, *Byte*, Vol 14, No 2, p 251-254.

Marshall (1989) Trevor and J.M. Tazelaar, Worth the RISC, *Byte*, Vol 14, No 2, pp 245-250.

Nicholls, William (1989) The Current Crop, *Byte*, Vol 14, No 2, p 235-245.

Norton, Peter (1986) *Inside the IBM PC*, Simon and Schuster, New York, pp. 387.

Poor, Alfred (1989) ADD-IN BOARDS - 16-bit VGA Cards Stretch the Standard, *PC Magazine*, Vol 8, No 13, 145-169.

Rosch, Winn L. and Kellyn S. Betts, Multiscanning Monitors for VGA and Beyond, *PC Magazine*, Vol 8, 9, (1989), pp 94-132.

Tanenbaum, Andrew S., Operating Systems - Design and Implementation, Prentice-Hall, Englewood Cliffs, NJ, 1987, pp.719.

The assessment of computer performance and computer speed for personal computers and desktop workstations is an often controversial subject. Since computers perform a wide variety of tasks in their applications, the results of benchmark evaluations are highly dependent on the tasks included in the test. There are several benchmarks used to evaluate computer performance, but unless the suite of tasks in the benchmark reflects the applications for your particular system, the results may not be relevant. This appendix will define some of the more common measures of computer speed that will be used to compare systems and define basic performance levels.

Clock Speed

The speed of the internal clock in the central processor determines the machine cycle time of the processor, which is the basic unit of measure for performing tasks. Each task the processor performs, such as getting an instruction from memory, interpreting the instruction, loading a register, storing data in memory, executing an addition, etc., takes one or more machine cycle times. Nominally, the faster the clock speed and the shorter the cycle time, the faster the computer. However, the architecture of the instruction set must also be included. On some chips the more complex instructions may take several cycles, thus slowing down the processor. Also, if memory access time or I/O tasks are slow, this will slow down the overall computing process no matter how fast the processor and clock speed. The ability of some processors to have more than one instruction in various stages of execution at a time, known as "pipelining", can increase overall performance even if the clock speed and cycle time are slow.

Instruction Speed

<u>MIPS - Millions of Instructions Per Second</u>: This is a measure of how many instructions can be executed in a second. However, this measure is not as straightforward as it seems. It is highly dependent on the instruction set used. The complexity and capability of the instruction set has a great impact on this measure. One class of chip may take five or six load and store instructions to move one piece of data from one storage location to another while another chip with a more complex instruction set may do it with one or two instructions. While the first chip may perform more instructions per second, it may actually take it longer to perform the load and store task. MIPS comparison should only be done for the same class of chips.

<u>MFLOPS - Millions of Floating-Point Operations Per Second</u>: Floating point calculations are more complex than the typically binary or integer type calculations that make up the basic instruction set used for MIPS determination. Again, the architecture of the chip plays a big role in performance. Some chips may be more efficient in single precision rather than double precision. For complete systems, the presence of a hardware co-processor that has hardware to do the floating point calculations rather than software algorithms will greatly impact speed.

Other Factors Affecting Performance

<u>Memory Access Speed</u>: The speed with which data can be accessed from memory by the CPU is a critical determinant of overall performance. The final access speed

involves two factors. First is the actual access time for a single byte of data. This is a function of the CPU design. The second factor is the bandwidth of information that can be passed at one time. A 32-bit memory board (like the Intel 80386) can pass 32 bits of information at one time. A 16-bit memory system would require two passes, while an 8-bit memory (like the 8088 chip in the original IBM PC) requires four passes.

I/O Access Speed: The speed with which data can be accessed from peripheral storage devices, such as hard disks, tape drives, and floppy disks, has an impact on system performance. As with memory access, the two factors determining access speed are the basic speed of the system and the bandwidth of the data transfer. The bandwidth is a function of the peripheral device and the bus or interface with the computer. A 32-bit bus can pass more data in a single step than an 8-bit bus.

Memory conflicts: Some peripherals, particularly video displays, also use memory for storage. This presents a conflict when the processor and the video display both need access to memory. This can result in the video "stealing" cycles from the processor and slowing down the computer. Most of the more powerful video display cards have their own memory to prevent this impact.

Software inefficiencies: Some application software may not take advantage of system performance capabilities, such as co-processors, complex instruction sets, and video memory, which degrades overall performance.

Peripheral and I/O System Conflicts: Waits for I/O processing and system and peripheral device interrupts break into the processing cycle to perform other tasks, which slows the processor down.

Baud: The rate of transmission of single piece of information per second, used in the form of *baud rate*. In applications where information is in the form of single bits, 1 Baud is the same as one-bit-per-second. However, some applications require more than one bit to transfer a piece of information (e.g. information may include check bits, parity bits, etc.), so Baud is not always equivalent to bits per second.

Bit: A single binary digit. It can have the value 0 or 1.

Bus: The bus is a set of parallel wires used to transmit data, commands, or power. For a data or command communications bus, the term is often used more generally to include the I/O control unit for a particular peripheral device that links the device to the computer bus and the channel device that links the processor to the bus.

Byte: Normally made up of 8 consecutive bits. On many computer systems it is the lowest addressable unit of main memory. It takes one byte of data to represent one ASCII character. The numeric values that can be represented by a byte range from 0 to 255 decimal.

Cache: Cache is a buffer in the CPU where data can be stored for ready access so the CPU does not have to go to memory to get data. For many computer chips, particularly RISC chips, the CPU instruction execution time is much faster than the memory access time. Constant memory access will slow the computer down. Most of the more sophisticated RISC chips use buffers or registers in the CPU as a cache to speed data access and increase performance. The same cache concept can also be used to increase I/O performance with hard disks. In this case, a cache buffer is created in memory that holds the most recently accessed data from the hard disk. Although memory access time is slower than access to buffers or registers in the CPU, memory access time is much faster than I/O access time for hard disks. Since most processes involve reading and writing the same blocks of data (for example, word processing and data base updates), storing data in memory rather than continuous reading and writing to the hard disk, increases performance. The drawback to using a cache is that it takes up memory space, which limits the amount of memory available for programs.

Central Processing Unit (CPU): The CPU or processor is the heart of a computer chip. It is the component that selects and executes the instructions for the computer. It typically includes a clock, an instruction control unit to select and read the instructions for execution, an arithmetic and logic unit to perform the instructions, and registers and buffers to hold data, addresses and the results of instructions. An integrated circuit, often referred to as a computer-on-a-chip, includes the basic CPU, main memory, plus other buffers registers and control units that support the complete computer system. The Intel 80286 and 80386 and the Motorola 68020 chips are examples of integrated circuits that contain CPUs.

Co-processor or Numeric Processing Extension (NPX): Usually refers to an auxiliary support chip to the CPU which is used to perform floating point calculations and special mathematical functions, such as sine, cosine, log and square root. The basic computer chip performs floating point calculations using software algorithms. By adding a co-processor chip that does these calculations via hardware,

the computing speed is increased significantly. For scientific calculations, a co-processor is usually necessary.

Networking: A means of linking two or more computer systems to share resources and to exchange data and information.

Network Server: A computer that donates resources to a network. Examples of functions a server can perform include: a communication hub that directs message traffic and keeps track of the addresses of stations on the network or a network poller that manages communication traffic (a network manager); a central-site resource for mass storage of files and data (file server); a main frame site to support complex computational problems and models (computation server); or a dedicated system with various peripheral devices that are shared by multiple-users.

Network Client: An individual workstation on the network. Each station has its own address. The client can use the network to transfer data to other stations, transfer data to and from the main frame server, run applications programs on the server, etc.

Ethernet: An IEEE standard originally developed by DARPA to allow high speed serial, asynchronous communications between computers. Ethernet is usually implemented in a bus topology consisting of well defined ends. Local nodes are tapped into the trunk line or bus. The standard defines the electrical signal specifications of the cable itself. The Ethernet system uses a technique known as collision detection to determine if a transmitted message has been interfered with by other messages on the network. Ethernet is rated at 10 Million bits per second or about 1 megabyte per second. Because of the collision detection process, the efficiency decreases as the frequency of use increases. The actual efficiency of its use is also dependent on the intelligence of the network server and on the speed of the computers. For workstation-class machines, the actual throughput is typically a few hundred kilobytes per second. The definitions for the actual information configuration, data format and operating procedures and protocols are not part of this standard and a protocol such as DECnet or TCP/IP must be used. Other IEEE standards exist for other network media and topologies. GL has an Ethernet network that is served by the VAX cluster using DECnet.

Token Ring: A network topology in which the media cable is configured in a closed loop. Local nodes on the cable are spliced into the cable as opposed to being tapped as in an Ethernet bus. The network server sends a `token' to each of the stations on the ring in-turn to put data or messages on the network.

Star Network: The Star Network is another topology for a network where each of the individual workstations is tied directly into the Network Server is a star fashion, as opposed the Ethernet Bus and the Token ring. For large, dispersed networks this topology is unwieldy and degraded by signal losses in long cables.

Transmission Control Protocol/Internet Protocol (TCP/IP): A public domain protocol developed by DARPA which works on a Ethernet network to transfer data between computers. Developed originally around machines using the UNIX operating system, this public domain protocol has been adapted by several vendors and is accessible through various operating systems. The Internet Protocol establishes the procedures for addressing other nodes on the network, including the paths to various gateways and nodes required to attain access to another address. The Transmission Control Protocol establishes the rules and procedures for coding,

decoding, and processing the data packets that are transmitted by the Ethernet network protocols.

DECnet: A proprietary software system developed by Digital Equipment Corporation (DEC) that implements the Ethernet network protocol. A separate hardware controller or interface handles the actual network interaction. DECnet software performs the same functions as TCP/IP software and drives the network interface and codes and decodes the transmitted data.

Multi-tasking (multiprogramming): The ability of the operating system to apparently execute two or more tasks at the same time. Multi-tasking operating systems actually multiplex the system resources between the current tasks several times per second thus giving the appearance of congruent execution. In fact the system is not processing two tasks at the same time. For example, while one task is writing data to a disk (an I/O function that doesn't require use of the CPU), another task can be doing calculations in the CPU.

Multi-user (time sharing): The ability of the operating system to treat simultaneous working sessions of different user accounts separately. Each user will have his or her own log in directory with sub-directories. The user is free to add, delete and move files and data within this environment and delegate different levels of authority to other users as he wishes. Usually multi-user operating systems also include multi-tasking.

Operating System (OS): The system level software that 1) manages the computer system resources (CPU, memory, disk drives, graphics terminals, etc.) on behalf of the user and his programs, and 2) provides a friendly interface between the user and the computer to accept commands, and 3) controls the programs that are run. The OS handles the awkwardness of the machine level language and the complexities of the hardware so they are transparent to the user. The architecture, capabilities, and limitations of the Operating System define the basic capabilities of the total computer system. For example, an Intel 80386 processor can support both a Microsoft Disk Operating System (single task DOS) and a Unix operating system (multi-tasking). Classes of operating systems include: single user-single task (for example, DOS), multi-tasking (for example, Unix), Virtual Memory multi-user operating systems (for example, DEC's VMS), etc.

Disk Operating System (DOS): The operating system tailored by Microsoft Corporation for the IBM PC. Because of the success of the IBM PC-type computers, it has become an industry standard.

VMS (Virtual Memory System): A proprietary operating system designed by the Digital Equipment Corporation (DEC) for implementation on the VAX line of computers. VMS is a virtual memory, multi-user, multi-tasking operating system.

UNIX: A virtual memory, multi-user, multi-tasking operating system originally developed by AT&T on the PDP line of DEC computers. It was originally developed as a document preparation and tracking operating system. UNIX was adopted by DARPA to be the standard operating system within the Department of Defense. UNIX is one of the few, if not only, operating systems designed to be transportable to several different hardware configurations. This transportability is largely the result of writing the

operating system in the high level language C. Thus, in principle, any machine that has a C compiler is capable of implementing the UNIX operating system.

Sun Operating System (SOS): A proprietary operating system designed by SUN Microsystems Inc. based on the UNIX operating system. By license agreement with AT&T, all UNIX based operating systems must provide a set of tools and constructs which existed in the original UNIX operating system. Thus the proprietary nature of this operating system is largely reduced to enhancements of the original operating system. SOS was designed to be portable to several hardware platforms including the Motorola 680XX, Intel 80386, and SPARC CPUs.

Ultrix: A proprietary operating system based on the UNIX operating system and built by Digital Equipment Corporation for implementation on the VAX line of computers. By license agreement with AT&T, all UNIX based operating systems must provide a set of tools and constructs that existed in the original UNIX operating system. Thus the proprietary nature of this operating system is largely reduced to enhancements of the original operating system. Ultrix is a virtual memory, multi-user, multi-tasking operating system.

Xenix: A Microsoft version of the UNIX operating system available commercially for use with personal computer systems.

Proprietary Standards: Standards for software or hardware that are controlled by license or patent by the developing organization. Examples are the majority of the DEC software and hardware, Silicon Graphics and Sun operating system for their graphics workstations, and most commercially available DOS-related applications software.

Public Domain Standards: A general classification of both hardware and usually software that has been produced or given to the public domain. Public domain means that all the specifications needed to implement the application can be found or obtained in the open literature. Some typical standards include: Ethernet, TCP/IP, NFS, Berkeley UNIX, FITS, PHIGS, MIT's X-Windows, PostScript, etc. The term open architecture is also often used to refer to public domain standards. The standards are usually approved by a public agency or professional organization, such as NITS, ANSI or IEEE.

Reduced Instruction Set Computer (RISC): A new breed of computer chip that uses a less complicated instruction set in its design. Although it may require more instructions to perform a particular function, the chip is designed so that one instruction is performed each clock cycle, so the result is a net increase in speed over the Complex Instruction Set Computer (CISC) chips, like the Intel 80386 or the Motorola 68020. Many of the RISC chips can actually perform two or more instructions at once, for example, instructions that involve main memory, RAM, the disk, and I/O devices can be done simultaneously. RISC chips also use a large number of registers to reduce memory access requirements and most use some form of cache buffer. Examples of new RISC chips include the Intel 80960, the Motorola 88100, the Advanced Micro Devices 29000, and the MIPS Computer Systems R2000. Sun Microsystems developed the Scalable Processor Architecture (SPARC) family of RISC chips that are made by a number of manufacturers.

Virtual Memory: Virtual memory is the relative memory space that a program takes in computer memory. When a program is designed, only virtual or relative memory addresses are used. Therefore, the program can be located anywhere in the physical memory and still operate. The actual allocation of physical memory for a program, that is, the assignment of specific memory locations with fixed addresses, is controlled by the memory management unit. Virtual memory was first developed to replace overlaying as a way to handle large programs that exceeded the available memory. Rather than have the programmer design the overlays, the computer uses virtual memory assignments to keep only those portions of the program in memory that were required at any particular time. The portions of the program that are not needed in memory at a particular time, or cannot fit, are "swapped out" to the hard disk for temporary storage and then "swapped in" to memory when they are needed. Virtual memory assignment also allows for more than one program to be in memory at the same time, which is essential for the implementation of multi-tasking and multi-user systems.

SUPPLEMENTARY

INFORMATION

AD-A219991

ERRATA

GL-TR-89-0174
INSTRUMENTATION PAPERS, NO. 338

The Zenith Z-248 as a Scientific Workstation

R. R. BABCOCK
R. E. McINERNEY
B. T. BANCROFT
R. J. RAISTRICK

1 September 1989

ERRATA

NOTE THE CORRECT IDENTIFICATION OF AUTHORS IN BLOCK 6 AND 11 ON THE ATTACHED REPORT DOCUMENTATION PAGE.

GEOPHYSICS LABORATORY
AIR FORCE SYSTEMS COMMAND
UNITED STATES AIR FORCE
HANSCOM AIR FORCE BASE, MASSACHUSETTS 01731-5000

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. AGENCY USE ONLY (*Leave blank*)

2. REPORT DATE
1 September 1989

3. REPORT TYPE AND DATES COVERED
Interim Scientific

4. TITLE AND SUBTITLE
The Zenith Z-248 as a Scientific Workstation

5. FUNDING NUMBERS
PE62101F
PR76590500

6. AUTHOR(S)
R.R. Babcock *
R.E. McInerney
B.T. Bancroft *
R.J. Raistrick

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
Geophysics Laboratory (LCY)
Hanscom AFB, MA 01731-5000

8. PERFORMING ORGANIZATION REPORT NUMBER
GL-TR-89-0174
IP No. 338

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

10. SPONSORING/MONITORING AGENCY REPORT NUMBER

11. SUPPLEMENTARY NOTES
* Mission Research Corp, One Tara Blvd, Suite 302, Nashua, NH 03062

12a. DISTRIBUTION/AVAILABILITY STATEMENT

Approved for public release, distribution unlimited

12b. DISTRIBUTION CODE

13. ABSTRACT (*Maximum 200 words)–(Continue on reverse if necessary and identify by block number)*

This report reviews the development of the desktop computer and workstation for scientific data analysis and develops the basic concepts of connectivity, processing power, high-resolution graphical displays, and large local data storage that support the creation of the scientific workstation approach to geophysical data analysis. The report reviews the capabilities of the basic Z-248 configuration and identifies key enhancements in memory expansion, high resolution displays, added storage, and networking capabilities that can turn the Z-248 into a scientific workstation.

14. SUBJECT TERMS
Zenith Z-248
Personal computer
Scientific workstation

15. NUMBER OF PAGES
52

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
Unclassified	Unclassified	Unclassified	SAR

NSN 7540-01-280-5500

Standard Form 298 (Rev. 2-89)
Prescribed by ANSI Std Z39-18
298-102